MODIS DATA STUDY
TEAM PRESENTATION

May 11, 1990

AGENDA

1. MODIS Tradeoff Study: Store or Recompute Geometric Parameters (McKay)
2. MODIS Processing and Storage Requirements Document Delivery (Ardanuy/Team)

MODIS Tradeoff Study: Store or Recompute Geometric Data?

1. Background and Motivation

In the process of estimating data storage requirements for MODIS data products, the MODIS Data Study Team has encountered several instances in which required data items referenced more than once during processing could either be initially computed, stored in long-term storage, and retrieved as the need arises or recomputed from stored basic input numbers when the need occurs. A noteable example of this situation is interpolation between earth-located "anchor points" to determine earth locations and observation geometry for individual instrument observations.

This document is a "first-cut" attempt to establish a methodology for addressing questions of this nature. This is a preliminary analysis inasmuch as the criterion used to determine the relative merit of alternative approaches is a simple comparison of the time required to complete the alternative procedures. Since CPU time on a high-capability data system is very likely much more valuable than I/O time for a peripheral device, this analysis does not reflect the true economics of the situations considered. This analysis was done as a preliminary exercise to facilitate the development of a final, more realistic model and because the relative speed results are interesting, in and of themselves.

2. Technical Approach and Assumptions

In this analysis it is presumed that earth-location and geometric parameters required for each instrument observation pixel have been precisely computed at an array of "anchor points" chosen for their suitability as interpolation references. Required quantities for the remaining pixels are to be determined using interpolation between these anchor points. The basic question addressed in this analysis is "Is it better to interpolate just once between the anchor points, store the results, and retrieve the stored quantities as needed, or is it better to store just the anchor point information and recompute the interpolation information each time the need arises?". Results will affect product storage requirements and processor requirements. Obviously, the answer depends on the speed of the I/O devices available for retrieving data from storage and the speed of the processor available for use in recomputing results.

The basic information assumed in the analysis of MODIS-N tradeoffs is given in Table 1. Corresponding information for MODIS-T is given in Tabel 2. For the analysis it is presumed that the various spectral channels are co-registered, i.e. pixel location information is the same for all specral channels corresponding to a given nominal location. The anchor point distribution chosen has been analyzed and is thought to be suitable for MODIS purposes. The specific variables needed and their assumed sizes are listed in the tables. Since the recompute option would leave the

Table 1.

Parameters used in tradeoff study:
store or recompute MODIS-N geometric parameters.

Number of frames/scan	1582
Number of pixels to be located/frame	8
Number of anchor points/scan line	94
Number of scan lines containing anchor points	2
Number of variables/anchor point or pixel	6

Definition of variables	Length (bytes)
Latitude	4
Longitude	4
Solar-Zenith Angle	2
Satellite-Zenith Angle	2
Relative Azimuth	2
True Azimuth	2
Total bytes/anchor point or pixel	16

Total bytes stored/scan (Store option)	202496
Total bytes stored/scan (Recompute option)	3008
Total FLOP to retrieve data/scan (Store option)	75936
Total FLOPS. to interpolate/scan (Recompute option) (see MDST handout January 26, 1990, item no. 9)	1470000
Slope	0.14 (MBytes/sec)/MFLOPS
Reciprocal factor	7 MFLOPS/(MBytes/sec)

Table 2.

Parameters used in tradeoff study:
store or recompute MODIS-T geometric parameters.

Number of frames/scan	1007
Number of pixel locations/frame	30
Number of anchor points/scan line	80
Number of scan lines containing anchor points	5
Number of variables/anchor point or pixel	6

Definition of variables	Length (bytes)
Latitude	4
Longitude	4
Solar-Zenith Angle	2
Satellite-Zenith Angle	2
Relative Azimuth	2
True Azimuth	2
Total bytes/anchor point or pixel	16

Total bytes stored/scan (Store option)	483360
Total bytes stored/scan (Recompute option)	6400
Total FLOP to retrieve data/scan (Store option)	181260
Total FLOP to interpolate/scan (Recompute option) (see MDST handout January 26, 1990, item no. 9)	3870000
Slope	0.13 (MBytes/sec)/MFLOPS
Reciprocal factor	8 MFLOPS/(MBytes/sec)

interpolation results available in the CPU registers of the processor, a nominal one instruction per data item processing requirement has been assigned to the storage-retrieval option. This would allow for the retrieval of stored data to corresponding CPU register locations. The processing requirement for the interpolations is derived from a detailed analysis of interpolation requirements presented earlier (Item no. 9, MODIS Data Study Team "weekly handout", January 26, 1990).

3. Results

Analysis results are expressed in terms of the speed of data retrieval from storage (MBytes/sec) and processor speed (MFLOPS). For each MODIS instrument, the space of these variable is divided into two regions (see Figure 1); above the line, variables can be recomputed more quickly than they can be retrieved from storage; below the line, retrieval from storage is quicker than recomputation. The line represents that combination of I/O transfer rate and processing speed at which time requirements for retrieval and recomputation are just equal. The slope of this line is listed in the tables as the "Reciprocal factor". Results basically say that the processor speed, expressed in MFLOPS, must at least equal 7 (MODIS-N) or 8 (MODIS-T) times the I/O rate, expressed in MBytes/sec, if recomputation is advantageous. Otherwise, storage and retrieval is advantageous.

Representative processor and I/O speeds are needed to interpret results. Table 3 is a list of representative computers and their processing capacities (Dongarra, Jack J., Performance of Various Computers Using Standard Linear Equations Software in a Fortran Environment, Technical Memorandum No. 23, Argonne National Laboratory, Argonne, IL, May 11, 1989). Data transfer rates for representative systems are listed in Table 4.

Optical tape (CREO)	3.0 MBytes/sec
On-line magnetic disk (Cray)	9.6
Magnetic tape (Cray)	4.5
Magneto-optical (Summus LightDisk 650)	680 KBytes/sec
WORM	304
CD-ROM	150

Processing Strategy for Level−1 Navigation

Storage of IFOV Locations Versus Recomputation from Anchor Points

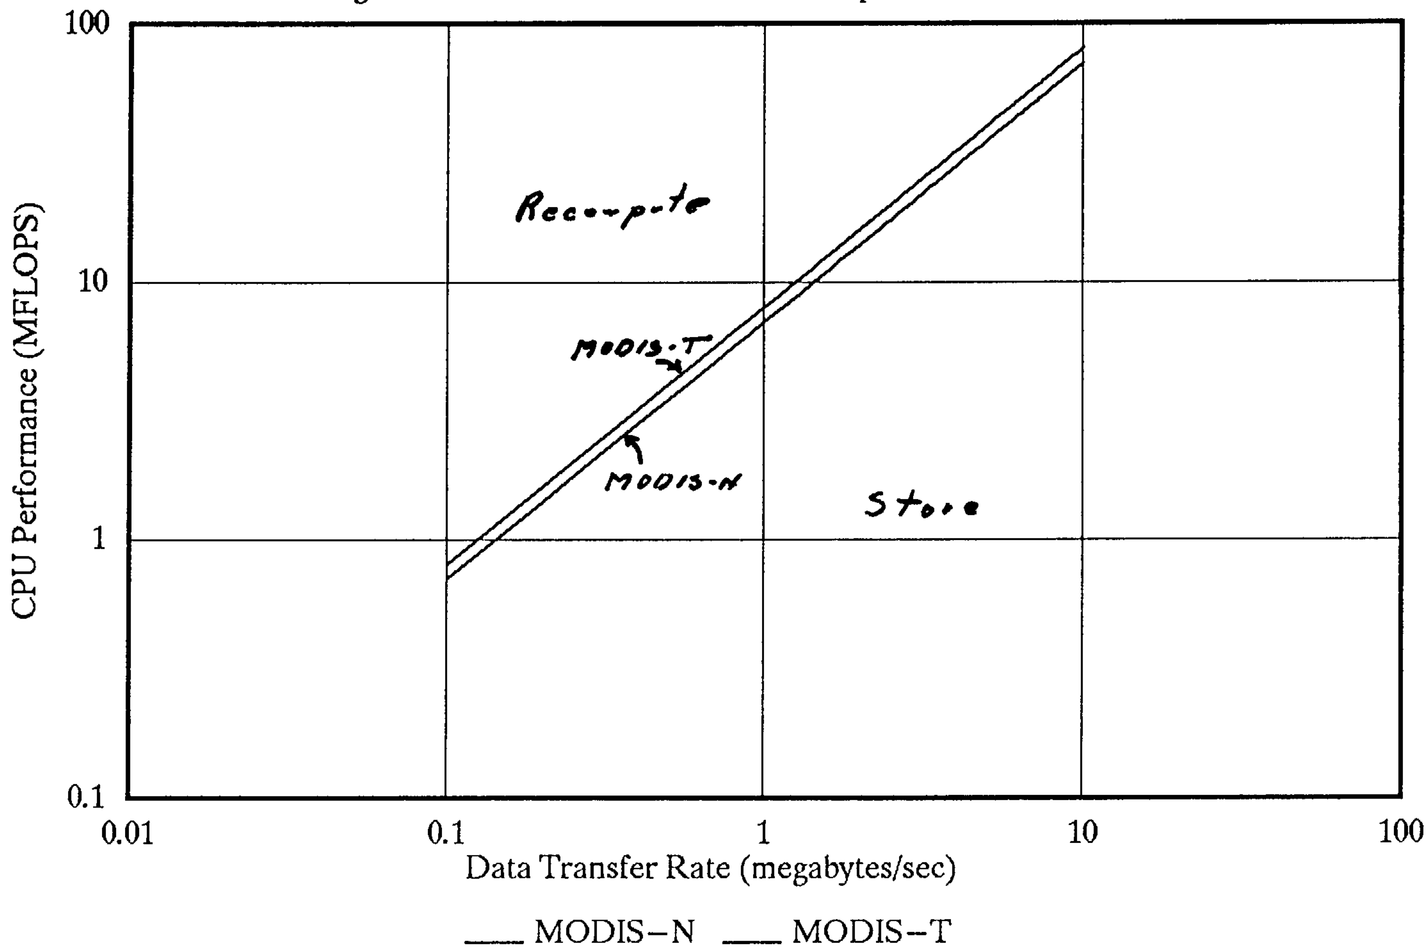

TABLE 1

Solving a System of Linear Equations

with LINPACK [a] in Full Precision [b] Using All Fortran

Computer	OS/Compiler[c]	Ratio[d]	MFLOPS[e]	Time secs
CRAY Y-MP/832 (8 proc. 6ns)	cf77 -Zp -Wm,-e78(Rolled BLAS)	.061	200	.00343
CRAY Y-MP/832 (4 proc. 6ns)	cf77 -Zp -Wm,-e78(Rolled BLAS)	.066	185	.00371
CRAY X-MP/416 (4 proc. 8.5ns)	cf77 3.0(Rolled BLAS)	.082	149	.00460
CRAY Y-MP/832 (2 proc. 6ns)	cf77 -Zp -Wm,-e78(Rolled BLAS)	.095	129	.00530
CRAY X-MP/416 (2 proc. 8.5ns)	cf77 3.0(Rolled BLAS)	.12	103	.00664
ETA10-G (1 proc. 7ns)	ETAV/FTN200(Rolled BLAS)	.13	93	.00736
CRAY Y-MP/832 (1 proc. 6ns)	cf77 -Zp -Wm,-e78(Rolled BLAS)	.15	84	.00814
CRAY 2S/4-128 (4 proc. 4.1ns)	cf77 3.0(Rolled BLAS)	.15	82	.00834
CRAY X-MP/416 (1 proc. 8.5ns)	cf77 3.0(Rolled BLAS)	.18	66	.0103
ETA10-E (1 proc. 10.5ns)	ETAV/FTN200(Rolled BLAS)	.20	62	.0110
CRAY-2/4-256 (4 proc. 4.1ns)	cf77 3.0(Rolled BLAS)	.20	62	.0112
CRAY 2S/4-128 (2 proc. 4.1ns)	cf77 3.0(Rolled BLAS)	.22	56	.0123
CRAY X-MP/14se (10 ns)	cf77 3.0(Rolled BLAS)	.23	53	.0130
CRAY-2/4-256 (2 proc. 4.1ns)	cf77 3.0(Rolled BLAS)	.26	48	.0144
NEC SX-2	FORTRAN 77/SX(Rolled BLAS)	.28	43	.0160
CRAY 2S/4-128 (1 proc. 4.1ns)	cf77 3.0(Rolled BLAS)	.30	41	.0169
CRAY-2/4-256 (1 proc. 4.1ns)	cf77 3.0(Rolled BLAS)	.32	38	.0178
NEC SX-1	FORTRAN 77/SX(Rolled BLAS)	.34	36	.0190
Hitachi S-820/80	HAP V21.00(Rolled BLAS)	.34	36	.0190
ETA10-Q (1 proc. 19 ns)	ETAV/FTN200(Rolled BLAS)	.36	34	.0200
NEC SX-1E	FORTRAN 77/SX(Rolled BLAS)	.38	32	.0210
ETA10-P (1 proc. 24ns)	ETAV/FTN200(Rolled BLAS)	.45	27	.0250
CRAY-1S	cf77 2.1(Rolled BLAS)	.46	27	.0260
CONVEX C-240 (4 proc)	Fortran 5.1(Rolled BLAS)	.48	25	.0270
NAS AS/EX 60 VPF	VAST/VS 1.4.1 opt=3	.52	24	.0289
CONVEX C-220 (2 proc)	Fortran 5.1(Rolled BLAS)	.60	21	.0324
Fujitsu VP-400	Fortran 77 V10L30(Rolled BLAS)	.61	20	.034
Amdahl 1400	77/VP V10L20(Rolled BLAS)	.64	19	.036
Amdahl 1200	77/VP V10L20(Rolled BLAS)	.65	19	.036
Fujitsu VP-200	Fortran 77(Rolled BLAS)	.70	18	.039
CDC Cyber 205 (2-pipe)	FTN(Rolled BLAS)	.70	17	.039
CONVEX C-210 (1 proc)	Fortran 5.1(Rolled BLAS)	.70	17	.039
Amdahl 1100	77/VP V10L20(Rolled BLAS)	.72	17	.040
Hitachi S-810/20	FORT77/HAP(Rolled BLAS)	.74	17	.042
Hitachi S-810/10	HAP V21.00(Rolled BLAS)	.75	16	.042
IBM 3090/180S VF	VS Fortran 2.3.0(Rolled BLAS)	.78	16	.043
Fujitsu VP-100	Fortran 77(Rolled BLAS)	.78	16	.044
Amdahl 500	77/VP V10L20(Rolled BLAS)	.78	16	.044
Siemens H120F	Fortran 77(Rolled BLAS)	.83	15	.0465
NAS AS/EX 60	VS 1.4.1 opt=3	.83	15	.0465
Cydrome Cydra 5	Fortran 77 Rel 2.4.1	.88	14	.049
Fujitsu VP-50	Fortran 77(Rolled BLAS)	.90	14	.051
IBM 3090/180E VF	VS 2.1.1 opt=3(Rolled BLAS)	.97	13	.054
CRAY-1S (1983 run)	**CFT 1.12 (Rolled BLAS)**	**1**	**12**	**.056**
IBM 3090/180 VF	VS Fortran V2(Rolled BLAS)	1.0	12	.057

Computer	OS/Compiler[c]	Ratio[d]	MFLOPS[e]	Time secs
CDC Cyber 990E	FTN V2 VL=HIGH	1.1	12	.059
Sperry 1100/90 ext w/ISP	UCS level 2	1.1	11	.061
Multiflow TRACE 14/200	Fortran 1.7(Rolled BLAS)	1.2	10	.068
Stellar GS 1000	f77 -O3 v2.0	1.3	9.8	.070
IBM 3090/180S	VS Fortran 2.3.0(Rolled BLAS)	1.3	9.6	.713
Ardent Titan-2 (2 proc)	f77 1.0 -O3 -inline(Rolled BLAS)	1.3	9.4	.073
IBM 3090/150E VF	VS 2.1.1 opt=3(Rolled BLAS)	1.4	8.8	.078
Alliant FX/80 (8 proc)	FX v4.0.21(RolledBLAS)	1.4	8.5	.0805
NAS AS/9160	VAST/VS 1.4.1 opt=3	1.5	8.3	.083
SCS-40	CFT 1.13(Rolled BLAS)	1.5	8.0	.086
IBM 3090/120E VF	VS 2.1.1 opt=3(Rolled BLAS)	1.6	7.5	.092
IBM 3090/180E	VS 2.1.1 opt=3	1.7	7.4	.093
CONVEX C-130	Fortran 4.0(Rolled BLAS)	1.7	7.3	.098
Siemens 7890F	Fortran 77 V10.3	1.7	7.2	.096
IBM 3090/180	VS opt=3	1.8	6.8	.102
Ardent Titan-1 (1 proc)	f77 1.0 -O2 -inline(Rolled BLAS)	1.89	6.5	.106
CONVEX C-120	fc 5.1(Rolled BLAS)	1.9	6.5	.106
Fujitsu M-380	Fortran 77, opt=3	1.9	6.3	.109
Multiflow TRACE 7/200	Fortran 1.4(Rolled BLAS)	2.0	6.0	.114
Alliant FX/80 (4 proc)	FX v4.0.21(RolledBLAS)	2.1	5.8	.119
Siemens 7890G	Fortran 77 V10.3 opt=4	2.1	5.9	.116
IBM 3090/150E	VS 2.1.1 opt=3	2.1	5.9	.117
FPS-264 (M64/60)	F02 APFTN64 OPT=4(Rolled BLAS)	2.1	5.9	.117
Alliant FX/40 (4 proc)	FX v4.0.21(RolledBLAS)	2.3	5.3	.129
Honeywell DPS90	ES F77V 1.0(Rolled BLAS)	2.5	5.0	.138
Siemens 7890D	Fortran 77 V10.3	2.4	5.0	.138
Alliant FX/4 (4 proc)	FX v4.0.21(RolledBLAS)	2.5	4.9	.139
CDC Cyber 875	FTN 5 opt=3	2.6	4.8	.143
CDC Cyber 176	FTN 5.1 opt=2	2.6	4.6	.148
NAS AS/EX 30	VS 1.4.1 opt=3	2.9	4.3	.162
Amdahl 5860 HSFPF	H enhanced opt=3	3.1	3.9	.176
Amdahl 5860 HSFPF	VS opt=3	3.2	3.8	.181
MIPS M/2000	f77	3.4	3.6	.191
NAS 8093 w/HSA	VS 1.4.0 opt=3	3.5	3.5	.197
Alliant FX/80 (2 proc)	FX v4.0.21(RolledBLAS)	3.6	3.4	.202
Alliant FX/40 (2 proc)	FX v4.0.21(RolledBLAS)	3.7	3.4	.205
CDC 7600	FTN	3.8	3.3	.210
Gould NP1	Fortran (Rolled BLAS)	3.9	3.1	.218
CDC Cyber 960-21	Fortran(Rolled BLAS)	3.9	3.1	.218
IBM 3090/120E	VS 2.1.1 opt=3	3.9	3.1	.221
CDC Cyber 960-31	NOS/VE 1.3.1 FTN 1.6	4.0	3.1	.222
FPS-264/20 (M64/50)	F02 APFTN64 OPT=4(Rolled BLAS)	4.1	3.0	.229
CONVEX C-1/XP	Fortran 2.0(Rolled BLAS)	4.1	3.0	.229
CONVEX C-1/XL	Fortran 1.6(Rolled BLAS)	4.2	2.9	.235
NAS AS/EX 25	VS 1.4.1 opt=3	4.2	2.9	.238
CDC Cyber 760	FTN 5, opt=3	4.7	2.6	.260
CyberPlus	CPFTN 1.1-07	4.8	2.6	.267
IBM 370/195	H enhanced opt=3	4.9	2.5	.275
NAS AS/EX 20	VS 1.4.1 opt=3	5.6	2.2	.312
IBM 3081 K (1 proc.)	H enhanced opt=3	5.7	2.1	.321
CDC Cyber 175	FTN 5 opt=2	5.8	2.1	.322